
# The Resource

U.S. ARMY ENGINEER RESEARCH AND DEVELOPMENT CENTER

**NEWSLETTER** 





maintaining the data needed, and coincluding suggestions for reducing	ection of information is estimated to ompleting and reviewing the collect this burden, to Washington Headqu ald be aware that notwithstanding and DMB control number.	ion of information. Send comments arters Services, Directorate for Info	s regarding this burden estimate ormation Operations and Reports	or any other aspect of the s, 1215 Jefferson Davis	his collection of information, Highway, Suite 1204, Arlington	
1. REPORT DATE <b>1999</b>		2. REPORT TYPE		3. DATES COVE 00-00-1999	ered <b>9 to 00-00-1999</b>	
4. TITLE AND SUBTITLE			5a. CONTRACT NUMBER			
The Resource. Win	ter 1999		5b. GRANT NUMBER			
					5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)		5d. PROJECT NUMBER				
					5e. TASK NUMBER	
					5f. WORK UNIT NUMBER	
U.S. Army Enginee	rce Center,3909 Ha	velopment Center,A	ATTN: ERDC	8. PERFORMING REPORT NUMB	G ORGANIZATION BER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)		
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAIL Approved for public	ABILITY STATEMENT	on unlimited				
13. SUPPLEMENTARY NO	TES					
14. ABSTRACT						
15. SUBJECT TERMS						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT unclassified	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>	Same as Report (SAR)	36		

**Report Documentation Page** 

Form Approved OMB No. 0704-0188

# From the Director's chair

October 1, 1999, was a memorable day for all of us at the U.S. Army Engineer Waterways Experiment Station (CEWES) in Vicksburg, MS. On that date, the 70-year-old research facility officially became a part of the newly organized U.S. Army Engineer Research and Development Center (ERDC). From a modest beginning as a single research facility initiated in response to the Great Flood of 1927, CEWES grew into a complex of five major engineering laboratories. With the most recent consolidation as ERDC, we are now eight laboratories located at four sites in the United States.

October 1, 1999, marked a new beginning for the CEWES MSRC as well, for on this day, we officially became the ERDC MSRC. Additional changes are inevitable, but without change there can be no progress. As with all memorable events, this is a good time to look back on our progress and successes and then look forward to the future — to revisit our vision and tighten our focus.

Beginning in 1989 as the Army Supercomputer Center, the CEWES MSRC became the first HPC Major Shared Resource Center in 1993 as part of the DoD HPC Modernization Program. Our growth has been phenomenal, and today we are ranked in the top ten in the world for computing capability.

We can measure our successes in numerous ways — in HPC computing resources (we house several of the most powerful HPC systems in the world), PET initiatives (we are capable of transferring cutting-edge HPC technology from premier university centers), scientific visualization capabilities (we have one of the most diverse and best-equipped SVC facilities of its kind), and our personnel (our staff consists of engineers, computer scientists, and visualization specialists) — the list goes on.



As we look to the future, we must continue to strive to excel and to deliver service, education, technical expertise, and leadership to the HPC community. Only then will we be able to achieve the research and engineering objectives that are vital to our nation.


The ERDC MSRC mission is to deliver HPC leadership, service, education, and technical expertise to achieve research and engineering objectives vital to the nation.

Bradley M. Comes Director, ERDC MSRC













## **Feature Articles**

A New Addition to the Family: The IBM Power3 SMP6
1999 Jackson State University HPC Summer Institute 8
Using Kerberos10
ERDC MSRC Common User Environment
— Overview14
Dr. Louis Turcotte to Serve as SC2000
Conference Chair16
SC99 Tutorial: Concurrent Programming
with Pthreads16
SC99 Poster Presentation: CH3D-SED16
SC99 Panel Discussion: Experiences with
Continuing OpenMP and MPI17
Perspectives on Scientific Visualization18
Scientific Visualization of the Future:
Getting a Grip on Large Data Sets18
Interview with Mike McCraney20
Interview with Richard Strelitz, Ph.D21
A Helping Hand in the Customer Assistance Center23
PET Profiles25
MSRC Regulars
Acronym List
Announcements 3
Off-campus4
Tool Tips
ERDC MSRC Training28
Technical Reports31
Training Schedule32

# acronym list

Below is a list of acronyms commonly used among the Department of Defense High Performance Computing community. You will find these acronyms throughout the articles in this newsletter.

AAU Administration, Allocation, and Utilization

API Application Programming Interface

AVS Application Visual Systems
CAC Customer Assistance Center
CDT Central Daylight Time

CFD Computational Fluid Dynamics

CHSSI Common High Performance Computing Software Support Initiative

CMG Computational Migration Group
COTS Commercial-off-the-Shelf Software
CSM Computational Structural Mechanics

C3I Command, Control, Communication, and Intelligence CWO Climate, Weather, and Ocean Modeling and Simulation

DoD Department of Defense DoE Department of Energy

EQM Environmental Quality Modeling and Simulation ERDC Engineer Research and Development Center

EWU Eastern Washington University FAQ Frequently Asked Questions

GB Gigabytes

GOTS Government-off-the-Shelf Software

HBCU/MI Historically Black Universities and Colleges/Minority Institutions

HPC High Performance Computing

HPCMO High Performance Computing Modernization Office HPCMP High Performance Computing Modernization Program

IT Information Technology

ITL Information Technology Laboratory

JSU Jackson State University

MIT Massachusetts Institute of Technology

MPI Message-Passing Interface MSF Mass Storage Facility

MSRC Major Shared Resource Center MSU Mississippi State University

NASA National Aeronautics and Space Administration

NSF National Science Foundation PBS Portable Batch System

PDPTA Parallel and Distributed Processing Techniques and Applications

PET Programming Environment and Training

PITAC Presidential Information Technology Advisory Committee

PL Performance Level PVM Parallel Virtual Machine

RF Radiofrequency

S/AAA Service Agency Approval Authorities

SC Supercomputing

SIAM Society for Industrial and Applied Mathematics

SMP Symmetric Multiprocessing SVC Scientific Visualization Center

SciVis Scientific Visualization

TICAM Texas Institute for Computational and Applied Mathematics

USACE U.S. Army Corps of Engineers

The contents of this newsletter are not to be used for advertising, publication, or promotional purposes. Citation of trade names does not constitute an official endorsement or approval or the use of such commercial products.

## announcements

### **ERDC: All systems go!**

Mary Gabb

It's official. The U.S. Army Engineer Waterways Experiment Station is now part of the U.S. Army Engineer Research and Development Center (ERDC). It may take a while for the local road signs to change but the transition at this site is complete. In an official ceremony on the WES front lawn on October 1, the ERDC was activated. Approximately 200 people attended the festivities, although the ceremony was broadcast live throughout the ERDC site in Vicksburg (ERDC command headquarters), as well as the ERDC laboratories in Champagne, IL, Hanover, NH, and Alexandria, VA.

MG Russell Fuhrman, Deputy Chief of Engineers of the U.S. Army Corps of Engineers (USACE), was the featured speaker at the ceremony. He expressed the significance of combining the eight USACE laboratories and the potential it offered to USACE, the Army, the DoD, and the nation into the 21st century. MG Fuhrman had served at the Vicksburg site during his 30-year career, and he remarked on the 70 years of tradition at Vicksburg. In his speech, he emphasized that the USACE system of laboratories gives the DoD the capability to push the technological envelope. It will be looked upon in the future as the single most important action we took to "stay ahead of the power curve...to stay efficient."

The ceremony also highlighted a time capsule in the new ERDC monument, which represents the combined strengths of the eight ERDC laboratories. Some of the time capsule's residents for the next several decades include a copy of the orders establishing ERDC, an organizational coin from each of the ERDC laborato-



MG Russell Fuhrman and COL Robin R. Cababa during the ERDC activation ceremony

ries, and a copy of the Corps' vision statement. Other momentos from the ceremony were also included, such as a disk containing a digital photo of the ceremony and a copy of the program.

The common theme of the ceremony, and indeed of the reorganization, is that the whole is greater than the sum of its parts. For ERDC, the collective strength of the eight laboratories will provide a state-of-the-art distributed team serving the numerous DoD technological interests.

A snapshot in time: The ERDC time capsule represents the synergy of the eight different research laboratories. Items in the capsule ranged from a digital photo of the ceremony to a copy of the orders establishing ERDC. The time capsule was placed in the ERDC monument unveiled on October 1, 1999.



# off-campus

### Customized Course on Parallel Programming with MPI and OpenMP

On July 6-8, 1999, a 3-day course entitled "Parallel Programming with MPI and OpenMP" was held at BBN Technologies (New London, CT) for ERDC MSRC users. The course was taught by Drs. Henry Gabb and Richard Weed of the ERDC MSRC Computational Migration Group and Programming Environment and Training Program, respectively. Scientists and engineers from BBN Technologies and Electric Boat Corporation (under contract to Navy) attended.

The course began with an overview of multiprocessor architectures and parallel computing. This was followed by a detailed discussion of OpenMP syntax and usage. A subset of functions from the MPI library was covered on the second day, along with a brief lecture on data layout for performance optimization. The last day of the course was devoted to OpenMP performance issues on the SGI Origin2000 as well as the benefits and pitfalls of combining MPI and OpenMP in the same program. (For more information on this course, please see page 28 ["ERDC MSRC Training"].)

### Invited Presentation at SIAM Conferences

Professor Mary F. Wheeler, ERDC MSRC PET Academic Lead for EQM, delivered an invited presentation at the jointly held Ninth SIAM Conference on Parallel Processing for Scientific Computing and Fifth SIAM Conference on Mathematical and Computational Issues in the Geosciences, sponsored by the Society for Industrial and Applied Mathematics (SIAM) in San Antonio on March 24, 1999.

The title of the presentation was "Synthetic Environments for Modeling Subsurface and Surface Flows." Professor Wheeler described a synthetic environment involving water quality and modeling in reservoirs, aquifers, bays, and estuaries. A computational environment that allows for multiple physical models, multiple discretizations and solvers, and flexibility in coupling of different physical models can greatly increase the productivity of scientists and engineers. Within the DoD, this technology is used for optimal design and operation of installation restoration, enhancement alternatives, development of short- and long-term strategies for integrated management in support of installation environmental quality, and stewardship and conservation of natural and cultural resources.

Professor Wheeler holds the Ernest and Virginia Cockrell Chair in engineering at The University of Texas at Austin and is a member of the National Academy of Engineering. She is the Director of the Center for Subsurface Modeling, a part of the Texas Institute for Computational and Applied Mathematics (TICAM). Professor Wheeler has taught four classes as part of the ERDC PET Training Program.

### Fortran Pthreads and Linear Algebra

Dr. Clay Breshears presented the paper "Application of Fortran Pthreads to Linear Algebra and Scientific Computing" at the 41st Cray User Group Conference held in Minneapolis, MN, on May 24-28, 1999. This paper reported on work done by Breshears and co-authors Drs. Mark Fahey and Henry Gabb to demonstrate the efficacy of the Fortran POSIX Threads programming interface developed by Breshears and Gabb. Two linear algebra problems (matrix multiplication and Gaussian elimination) and two algorithms from the U.S. Air Force Command, Control, Communication, and Intelligence (C3I) Benchmark were used in the study. In each case, the execution time of the threaded codes across multiple processors was lower than the original serial versions.

# off-campus

### **ERDC MSRC PET Academic Leads on President's IT Committee**

Professors Joe Thompson, Larry Smarr, and Ken Kennedy, PET Academic Leads at the ERDC MSRC, serve on President Clinton's Presidential Information Technology Advisory Committee (PITAC). PITAC provides the President, the Office of Science and Technology Policy, and Federal agencies involved in Computing, Information, and Communications research and development with guidance on all areas of high performance computing, communications, and information technologies. Representing the research, education, and library communities and including network providers and representatives from critical industries, the Committee helps to guide the Administration's efforts to accelerate development and adoption of information technologies vital for American prosperity in the 21st century.

The PITAC report released in February 1999 (www.ccic.gov/ac/report/) provided the basis for the President's \$366M Information Technology for the Twenty-First Century (IT2) Initiative (www.ccic.gov/it2/) proposed to Congress. Professor Thompson is currently co-chairing the PITAC subcommittee that is evaluating the agency IT2 Implementation Plan.

Professor Thompson is a William L. Giles Distinguished Professor at Mississippi State University. Professor Smarr is Director of the National Computational Science Alliance and Professor of Physics and Astronomy at the University of Illinois at Urbana-Champaign. Professor Kennedy is Director of the Center for Research on Parallel Computation and Ann and John Doerr Professor of Computer Science at Rice University.

### Taking MPI\_Connect One Step Further

Dr. Clay Breshears presented a paper at the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'99), held June 28 - July 1, 1999. The focus of work was on the use of MPI\_Connect to run on three different, geographically separated systems; the total execution time of the test case (CGWAVE) was reduced from 3,112 min (over 2 days) to 12 min.

#### **References:**

- 1. Breshears, C. P., Bova, S. W., Cuicchi, C., Demirbilek, Z., and Gabb, H. A. (1999). "Using MPI\_Connect to distribute parallel applications across multiple platforms." *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '99)* 1, 383-389, CSREA Press.
- 2. Bova, S. W., Breshears, C. P., Cuicchi, C., Demirbilek, Z., and Gabb, H. A. "Dual-level parallel analysis of harbor wave response using MPI and OpenMP." *International Journal of High Performance Computing Applications*. In press.

# A New Addition to the Family: The IBM Power3 SMP

Henry A. Gabb, Ph.D., Director of Scientific Computing

The IBM Power3 SMP is due to arrive at the ERDC MSRC before the end of the calendar year. I am excited about the possibilities it will offer our users.

The IBM Power3 SMP is a distributed, sharedmemory system. In total, it will have 512 processors. Our Challenge and CHSSI users, in particular, who already use Message-Passing Interface (MPI) will see immediate—and substantial—benefits as their codes will be easily portable to the new system. Adding threads to their MPI programs should be straightforward. Challenge projects typically use MPI because they require maximum performance and scalability. Those users ask the big questions that require the biggest computers. Our long-term goal, however, is to introduce all ERDC MSRC users to dual-level parallelism, especially users who are unnecessarily limiting their problem sizes. Their projects (and their scientific questions) could expand significantly with assistance from ERDC MSRC staff. We would like to take non-Challenge users and turn them into potential Challenge users by removing artificial barriers to problem size. Many of our users, especially those migrating from the Cray C90, still think that 8 GB is a lot of memory. It isn't. The ERDC MSRC parallel systems have tens or hundreds of gigabytes of memory.

Why the IBM Power3 SMP is architecturally interesting

The IBM Power3 SMP uses a symmetric multiprocessor node system architecture. In other words, within a node, multiple processors are all connected to the same memory.

The IBM Power3 SMP being delivered to the ERDC MSRC has 64 8-processor SMPs clustered together. The nodes are connected by a high-speed network. This type of architecture is gaining prominence because additional processors can be added more easily and cheaply as an SMP node.



Advances in high performance computing happen fast. Our award winning work in the SC98 Challenge Demo ("Dual-level parallel analysis of harbor wave response using MPI and OpenMP") was considered novel just last year. <sup>1,2</sup> At that time, only a few projects combined message passing and threads in the same application. On the IBM Power3 SMP, it will be standard practice.

We will be offering several tools for our users who will be running on the IBM Power3 SMP:

- ➤ The KAP/Pro Toolset (Kuck and Associates, Inc.). This will be available for Fortran, C, and C++ to help users debug and optimize OpenMP.
- ➤ VAMPIR for optimizing communication in message-passing applications.
- ➤ Ensight for interactive postprocessing for visualization.

### References:

- 1. Bova, S. W., Breshears, C. P., Gabb, H. A., Eigenmann, R., Gaertner, G., Kuhn, B., and Salvini, S. "Parallel programming with message passing and directives." *SIAM News*. In press.
- 2. Bova, S. W., Breshears, C. P., Cuicchi, C., Demirbilek, Z., and Gabb, H. A. "Dual-level parallel analysis of harbor wave response using MPI and OpenMP." *International Journal of High Performance Computing Applications*. In press.

"The advantage of the IBM Power3 SMP is that you have two levels of parallelism, two memory hierarchies. And it's highly scalable."

- Dr. Henry A. Gabb



### More Information

VAMPIR: see article on page 12 TotalView: see article on page 13

## 1999 Jackson State University HPC Summer Institute

Wayne Mastin, Ph.D.



Drs. Alan Shih (University of Illinois) and Richard Strelitz (ERDC MSRC) offer an introduction to Scientific Visualization for HPC to the JSU students. Dr. Strelitz is the Acting Manager of the ERDC MSRC Scientific Visualization Center.

The third annual Jackson State University (JSU) HPC Summer Institute was held on June 14-25, 1999. This year's institute attracted 17 students from 7 historically black colleges and universities and minority institutions (HBCU/MIs), from Pennsylvania to Oklahoma.

During the 3 years it has been in place, the JSU HPC Summer Institute has introduced HPC to 53 minority students representing 10 HBCU/MIs. The institute is organized each year by JSU under the direction of Professor Willie Brown, the PET HBCU/MI Lead. The objective of the institute is to introduce a group of students from HBCU/MIs to the need for HPC in solving problems of national significance. This year's institute began with a tour of the ERDC in Vicksburg, MS, on June 14. The ERDC MSRC PET academic and onsite leads conducted daily lectures and laboratory sessions to introduce the students to the tools and techniques of HPC. The presentations covered all

computational technology and related areas supported by the ERDC MSRC. For example:

- ➢ Dr. Phu Luong, ERDC MSRC onsite EQM Lead for The University of Texas at Austin, talked about applications of multiblock grid generation techniques to coastal ocean modeling.
- Professor Mary Wheeler, the Texas Institute for Computational and Applied Mathematics (TICAM), The University of Texas at Austin, gave lectures on subsurface and surface water modeling. Graduate students Sharon Lozano and Jennifer Proft offered tutoring in parallel computation.
- Dr. David Littlefield, the Institute for Advanced Technology, The University of Texas at Austin, and Dr. Richard Weed, Mississippi State University (MSU) and the PET onsite Lead for CSM, conducted a day of lectures and laboratory exercises on structural mechanics and the finite

element structural analysis method. Dr. Littlefield presented the fundamental concepts and physical principles of structural mechanics in a morning lecture session. Dr. Weed conducted an afternoon laboratory session in which students used two different finite element programs to perform analyses on different types of example structural systems.

Professor Bharat Soni (MSU), Dr. Steve Bova (now at Sandia National Laboratories, NM, formerly the ERDC MSRC PET CFD onsite Lead for MSU), and Professor Roy Koomullil (MSU) presented an overview of CFD technology and its use in the HPC industry. Professor Soni explained the concepts of model discretization and grid generation and promoted the need for high performance computing. Professor Koomullil demonstrated grid generation and solution postprocessing software, while Dr. Bova used the "Java Virtual Wind Tunnel" from the Massachusetts Institute of Technology (MIT) to illustrate basic concepts in CFD.

The institute concluded with a closing ceremony at JSU on June 25. Each student who completed the institute was presented a certificate by Mr. Bradley Comes, ERDC MSRC Director.

### The JSU HPC Summer Institute Line-up

The following schedule lists the topics and presenters at this year's institute:

Tuesday, June 15, 1999: Climate/Weather/Ocean Modeling and Simulation Dr. David Welsh, Ohio State University Professor P. Sadayappan, Ohio State University Dr. Stephen Wornom, ERDC MSRC

Wednesday, June 16, 1999: Environmental Quality Modeling and Simulation Professor Mary Wheeler, The University of Texas at Austin Professor Clint Walker, The University of Texas at Austin Ms. Sharon Lozano, The University of Texas at Austin Ms. Jennifer Proft, The University of Texas at Austin Dr. Phu Luong, ERDC MSRC

Thursday, June 17, 1999: Computational Structural Mechanics Dr. David Littlefield, The University of Texas at Austin Dr. Richard Weed, ERDC MSRC

Friday, June 18, 1999: Computational Fluid Dynamics Professor Bharat Soni, Mississippi State University Dr. Steve Bova, ERDC MSRC Professor Roy Koomullil, MSU

Monday, June 21, 1999: Communication/Collaboration Professor Geoffrey Fox, Syracuse University Professor Nancy McCracken, Syracuse University

Tuesday, June 22, 1999: Scalable Parallel Programming Tools Dr. Graham Fagg, University of Tennessee, Knoxville Dr. Richard Hanson, Rice University

Wednesday, June 23, 1999: Scientific Visualization Dr. Alan Shih, University of Illinois Dr. Richard Strelitz, ERDC MSRC

Thursday, June 24, 1999: Forces Modeling and Simulation/C4I Professor Wojtek Furmanski, Syracuse University Mr. Tom Pulikal, Syracuse University

## **Using Kerberos**

Jay Cliburn

Effective October 1, 1999, we gate was renamed to agate. All references in this article are, therefore, as agate.

"Kerberos is a network authentication protocol. It is designed to provide strong authentication for client/server applications by using secret-key cryptography."

-The Kerberos FAQ

By now, virtually all ERDC MSRC customers are acquainted with the Kerberos security mechanism used to authenticate user identity on our computing resources. The most frequently used commands that employ Kerberos are **telnet**, **rlogin**, **ssh** (secure shell), and **ftp**. We find that many customers still access the ERDC MSRC through the bastion host, called agate, when in fact customers can access MSRC resources directly without passing through agate for authentication. This article discusses direct access to ERDC MSRC resources.

The critical prerequisite to directly accessing any ERDC MSRC system except agate is the user's possession of an active Kerberos ticket. A ticket is obtained by executing the **kinit** command on Unix machines or the **krb5.exe** command on non-Unix hosts. These commands are included in the Kerberos client distribution kit. If the machine you are using does not have the distribution kit installed because it is not available, then you must access ERDC MSRC resources through agate. Agate's *raison d'être* is to provide a means to acquire a Kerberos ticket and gain access to MSRC hosts when your own workstation does not contain the Kerberos client applications.

Beginning October 1, 1999, the ERDC MSRC changed the access policy for agate. The system will no longer accept non-Kerberized telnet access. This change was made to reduce the risk of sending the user's Kerberos password in the clear over the Internet. The system will still accept non-Kerberized **ssh** with SecurID authentication, but the user must submit a special access request to the Customer Assistance Center.

#### Unix users

If your workstation has the Kerberos client kit installed, then you do not and should not need to log in to agate to access ERDC MSRC hosts. From a Unix host, a typical direct-access session might look like this (words in bold font indicate text typed in by the user; words in regular font show the screen text):

```
yourhost$ /usr/local/bin/kinit
Password for your_principal@WES.HPC.MIL:
Passcode:
yourhost$ /usr/local/bin/telnet origin.wes.hpc.mil
Trying 134.164.13.43...
Connected to origin.wes.hpc.mil (134.164.13.43).
Escape character is `^]'.
[ Kerberos V5 accepts you as
"your_principal@WES.HPC.MIL" ]
[ Kerberos V5 accepted forwarded credentials ]
...
...
origin$
```

Note in the session above that the Kerberos commands are located in /usr/local/bin. This may not be the location of the commands on your particular workstation. Further, by default the client kit installs the commands as **ktelnet**, **kftp**, and so on instead of simply **telnet**, **ftp**, etc. If the complete Kerberos distribution has been installed by your system administrator, the standard commands, **telnet**, **ftp**, etc., may be the Kerberized versions.

In order for direct access to work, you must first acquire a Kerberos ticket, as was done with the **kinit** command above. You must also have a correctly formatted /etc/krb5.conf file installed. This file is included with the client distribution kit, but may need to be modified for your specific installation. Particularly, your *default realm* must be set to the site that issued your SecurID card. This is called your "home site." An improperly configured krb5.conf file is a common source of difficulty in accessing MSRC hosts.

If you wish to send an X-Window display from an MSRC host back to your local workstation, access the MSRC host with **ssh** instead of **telnet** or **rlogin. ssh** automatically sets up the proper security and display parameters such that the X-Window can be exported to and displayed on your workstation securely. *Never use* **xhost** to enable remote displays on your workstation. Using the **xhost** command can expose the whole X-session to hijacking.

#### WINTEL Users

If you use a Windows/Intel platform to access ERDC MSRC systems, your access options are somewhat more limited than those of Unix users. The Kerberos client distribution kit contains a single interactive access client: telnet.

To directly access MSRC hosts using a WINTEL machine, you must first obtain a ticket by executing the **krb5.exe** command. The Kerberos window will pop up, prompting you for a password. After entering your password, the KerbNet window will pop up, challenging you for your SecurID Passcode. If your password and passcode are correct, you will be granted a Kerberos ticket. If you are unable to obtain a ticket, the problem might lie in your krb5.ini file. The *default realm* entry in that file must be set to your home site (the site that issued your SecurID card). An improperly configured krb5.ini file is a common source of difficulty for WINTEL customers when accessing MSRC hosts.

Note that the access procedure of the preceding paragraph did not include logging in to agate. Access to agate is required only if you are unable to acquire a Kerberos ticket on your workstation, or if you wish to utilize a connection client other than telnet on a WINTEL workstation to gain access to ERDC MSRC hosts. For FAQ on Kerberos, please see <a href="http://www.nrl.navy.mil/CCS/help/FAQ.html">http://www.nrl.navy.mil/CCS/help/FAQ.html</a>.

### Conclusion

A number of ERDC MSRC customers who remotely access MSRC computing resources still obtain that access through the bastion host, agate. Under most circumstances, this intermediate step is unnecessary, and users should be able to gain direct access to MSRC systems. If you experience problems establishing a login session without agate in the path, please contact the ERDC MSRC Customer Assistance Center at *info-hpc@wes.hpc.mil* or call 1-800-500-4722.

### More information

List of supported client kits:

www.nrl.navy.mil/CCS/help/kerberos.support.nontable.html
Kerberos client kits are available via anonymous ftp from:
(You must read the README file and agree with its contents to access the final directory containing the client distributions.):

ftp://ftp.cmf.nrl.navy.mil/pub/kerberos5

Current krb5.cont and krb5.ini used by the MSRC can be obtained via anonymous ftp from: ftp://ftp.cmf.nrl.navy.mil/pub/kerberos5/krb5.conf ftp://ftp.cmf.nrl.navy.mil/pub/kerberos5/krb5.ini

Note: This software is export controlled and cannot be legally distributed outside the United States.

# tool tips


Clay P. Breshears, Ph.D.

### **VAMPIR**

VAMPIR is a visualization tool for analyzing the performance of message-passing programs. VAMPIR is able to graphically display trace data generated by VAMPIRtrace and saved in tracefiles generated from running instrumented code. In order to generate these tracefiles, the VAMPIRtrace Message-Passing Interface (MPI) profiling library is linked with a user's code to instrument all calls to MPI routines within the code. When executed, the instrumented MPI calls generate trace data detailing when a call was initiated and how long the function call lasted.

The VAMPIRtrace library can be used to instrument codes running on any of the ERDC MSRC HPC platforms. The VAMPIR visualization tool is available on all ERDC MSRC systems except the Cray T3E (SGI). Thus, tracefiles can be generated by MPI programs running on any HPC platform, but those generated on the T3E must be moved to another platform for visualization.

VAMPIR offers three types of visualization: timeline, statistics, and the process state displays. The timeline display shows process states over time and communication between processes by drawing lines to connect the message sender and receiver. Message patterns and time spent waiting for messages, completion of global communication routines, or other facets of execution can be easily seen with this display. Timeline display is good for finding communication bottlenecks. The statistics display shows the cumulative statistics for the complete tracefile in pie-chart form for each process. The percentage of execution time taken up by all communication or one particular MPI routine can be determined from this display. The process state display represents every process as a box and shows the process state at a selected point within the execution trace. This display allows the user to identify how many processes are executing MPI calls or user code, or standing idle. This display is good for identifying load imbalances, which adversely impact performance. Details and features of each of these displays



VAMPIR and VAMPIRtrace provide three main displays: timeline, statistics, and process state

can be configured by the user with pull-down menu options. Such customizations can be saved to a configuration file that is read each time VAMPIR is started.

The VAMPIRtrace instrumentation automatically gathers performance data for each MPI call in a program. In long-running applications with many MPI calls, a user may be interested in only a particular portion of the program. Tracing may be enabled and disabled at the discretion of the user by inserting calls to VAMPIRtrace Application Programming Interface (API) functions. In this way, performance analysis can be concentrated on specific areas of a code. This helps to control the amount of trace data generated, which can become quite large for programs that use many processors or execute for a long time. Also, because any parts of the code that are not MPI routines are considered to be part of the generic "User Code," the API provides calls for defining, starting, and stopping user-defined activities. Users can gather performance data on specific user-written routines in addition to the default MPI activity.

VAMPIR includes rapid zooming capabilities, forward/backward motion within the timeline display, and flexible filter operations to reduce the amount of information displayed. These options allow the user to focus quickly on arbitrary time intervals and easily identify performance bottlenecks at an appropriate level of detail.

#### **TotalView**

TotalView is a source-level, multiprocess debugger; it is extremely powerful and flexible. TotalView supports source-level debugging in C, C++, FORTRAN77, Fortran 90, and High-Performance Fortran. Because TotalView is thread-aware, it can debug OpenMP and Pthreads programs. It is also MPI-aware. It has a graphical user interface (based on the X-Windows System), which enables users to become productive with this tool in a short time. TotalView is available on all ERDC MSRC HPC platforms except on our latest addition, the IBM Power3 SMP (please see article on page 6).

TotalView can be used to either monitor a running program or perform a postmortem analysis on a core file. It offers a full range of functions for process control, including attaching to existing processes. The programmer can set, delete, enable, and disable breakpoints, conditional breakpoints, evaluation points, and event points at both the source code and machine code levels. These points enable the programmer to sequentially step through the code on a statement-by-statement basis, if necessary, providing "snapshots" of data at each cycle.

TotalView was designed for debugging multiprocess codes. Each process is presented in a separate window, each displaying information specific to that process. It is not necessary to display all process windows at the same time. The programmer may close any windows that are not of interest and focus debugging efforts on the relevant processes of the computation. Besides giving the user access to values stored in program variables, TotalView is able to graphically depict arrays of data during debugging sessions. This feature can be used to examine large amounts of data in a concise manner, locate incorrect values quickly, and watch data trends during program execution. Unlike textual views of data that are updated automatically, graphic visualizations must be updated by the user.

More Information
Customer Assistance Center:
info-hpc@wes.hpc.mil
www.wes.hpc.mil
1-800-500-4722

### **ERDC MSRC Common User Environment**

### — Overview

Jay Cliburn and Paul Adams

The ERDC MSRC is composed of multiple architectures, including systems from SGI/Cray (Unicos), SGI (IRIX), and IBM (AIX). Each architecture brings with it a set of operating environment constructs and elements of "look and feel" that have evolved to become distinct and vendor-specific over time. Because a single MSRC user may operate on more than one of these architectures, a variety of differences between the operating environment of each machine are often encountered. The ERDC MSRC has embarked upon an initiative to establish a common user environment to facilitate a smooth transition for users from one platform to another.

Key elements of the initiative include the following:

- ➤ A common batch queuing system
- Consistent application locations and directory names
- ➤ Identical group names (same typographical case)
- ➤ Common working area directory names
- Common bulk data transfer utilities.

Each of these elements is discussed below. Please note that this initiative is a work-in-progress and that any of these elements are subject to enhancement upon review and recommendation by the Common User Environment Working Group.

### **Common Batch Queuing System**

The Portable Batch System (PBS) has been selected as the batch queuing system for the ERDC MSRC. Currently PBS operates on the two IBM SPs and the SGI/Cray Origin2000 and will be functional on the Cray T3E by October 1. PBS is found in the /pbs directory on the SPs and the Origin2000. When installed on the T3E, it will be located in /pbs as well.

### **Consistent Application Locations and Directory Names**

The ERDC MSRC has chosen the following common directory structure for application and user software:

/usr/local/applic/	COTS/GOTS/supported application software
/usr/local/bin/	COTS/GOTS/other executables
/usr/local/lib/	COTS/GOTS/other libraries
/usr/local/doc	documentation
/usr/local/man	man pages
/usr/local/usp	unsupported user software

An exception to the above format exists on the IBM SPs, where application software is located in the /gpfs/cots directory tree. It is planned that this condition will be rectified through the use of symbolic links to /usr/local or by some other means.

### **Identical Group Names**

At the current time, certain groups are named in uppercase characters on some systems and all lowercase on others. For example, a member of the PET group might have a group name of PET (uppercase) on one system and pet (lowercase) on another. Upon completion of this initiative, all group names will be lowercase. Please note that this is a simple change to be done in the near future; it will have no impact on the user.

### **Common Working Area Directory Names**

"Common working area names" is actually something of a misnomer. The goal is not to have all work area directories across all systems carry the same name. Instead, a single environment variable should be used to access the work area, which allows the underlying directory or file system to carry any name. ERDC MSRC has chosen

\$WORKDIR as the variable name. This is the area to which input data decks are staged for queued and running jobs and to which output data are written before being transferred to mass storage. It is temporary storage. The following table illustrates the underlying directory names for \$WORKDIR:

Host \$WORKDIR
Origin2000 /Work/username
IBM SP (both) /gpfs/work/username
Cray T3E /tmp/username

The IBM SPs have an additional area of workspace. Each compute node contains a certain amount of local disk space, which the ERDC MSRC calls \$SCRATCH. \$WORKDIR is common to all compute nodes, while \$SCRATCH is visible only to each node to which the disk is locally mounted. The underlying directory name for \$SCRATCH is /scratch/username.

#### Common Bulk Data Transfer Utilities

The ERDC MSRC seeks to provide users with identical bulk data transfer utilities across all HPC hosts. These utilities are used to move data to and from the mass storage facility (MSF) to the host computer. The utilities employ remote file transfer protocol as the data pump and are listed here:

**msfget** Move a file from MSF to the current

login host

**msfput** Move a file from the current login host

to MSF

**msfmget** Move multiple files from MSF to the

current login host

**msfmput** Move multiple files from the current

login host to MSF

msfstat Check the availability of both MSF and

the user's home directory on MSF.

#### Conclusion


In an attempt to smooth the transition for users from one HPC platform to another, the ERDC MSRC has embarked upon the mission of establishing a common user environment. Because each machine has its own operating environment constructs, instances may occur where a common user environment is not possible. However, the goal is to provide a common batch queuing system, consistent application locations and directory names, identical group names (same typographical case), common working area directory names, and common bulk data transfer utilities. As of October 1, the task of providing users a common environment in which to operate is 95 percent complete at the ERDC MSRC.

More information

**Documentation/User Information:** www.wes.hpc.mil/documentation/index.html



# Dr. Louis Turcotte to Serve as SC2000 Conference Chair



volunteers and a conference budget of over

Dr. Turcotte has been involved in high performance computing for over 20 years. He has authored numerous articles, co-authored two books on computational methods, taught tutorials at SC, and frequently gives invited talks in topics related to high performance computing. His role in ITL, analogous to a chief technologist, includes programmatic involvement in technical activities throughout the laboratory. Among his various programmatic responsibilities, he serves as the PET Director for the ERDC MSRC, where he provides leadership in shaping the contributions of the 10 university partners.

# SC99 Tutorial: Concurrent Programming with Pthreads



Henry A. Gabb, Ph.D.

\$2,500,000.

Drs. Henry A. Gabb and Clay P. Breshears have been selected to conduct a tutorial on Pthreads at SC99 (Portland, OR; November 13-19, 1999). Pthreads is the POSIX standard library for multithreading in which tasks are assigned to threads that execute concurrently. On symmetric multiprocessors, threaded tasks can execute in parallel. The Pthreads library consists of 61 functions governing thread creation and management, mutual exclusion, condition variables and thread signaling, and low-level scheduling.

The tutorial begins with a discussion of concurrency. Classic concurrent programming models (e.g., bossworker, producer-consumer) and problems (e.g., The Dining Philosophers) illustrate the use of threads to express concurrent tasks. The pitfalls of race conditions and deadlock are introduced.

After laying a foundation in concurrency, the tutorial covers the 14 core Pthreads functions most useful to scientific programming. Each function is discussed in detail with example codes to illustrate usage. A detailed discussion of mutual exclusion variables shows how to avoid common race conditions like write/write and read/write conflicts. Lastly, condition variables and thread signaling are discussed.

Finally, a series of examples will be presented to demonstrate the utility of Pthreads in scientific computing. These include algorithms from the C3I Benchmark (terrain masking, route optimization, mapimage correlation) as well as matrix multiplication and LU decomposition.

### SC99 Poster Presentations: CH3D-SED

Drs. Phu Luong (PET EQM Lead), Clay P. Breshears (PET Scalable Parallel Programming Tools Lead), and Henry A. Gabb (Director of Scientific Computing) will present a poster entitled "Dual-level parallelism improves load balance in the production engineering application CH3D-SED" at SC99.

CH3D-SED is a three-dimensional, numerical, hydrodynamic, salinity, and temperature model. It is a production application widely used to investigate important physical features of the hydrodynamic process and bathymetry in areas such as the Chesapeake Bay, Delaware Bay, and New York Bight. These areas are frequently traveled by U.S. Navy vessels. Over the years, parallelization of CH3D-SED has been conducted. A parallel version with one-dimensional (1-D) domain decomposition has been used by many environmental scientists. Luong et al. will describe a dual-level parallelism implementation of the 1-D domain decomposition version of CH3D-SED. MPI is used to parallelize the domain decomposition. However, the workload varies between MPI processes. To improve load balance and performance, OpenMP is used to dynamically thread the computational domain. Performance results for MPI-only, OpenMP-only, and MPI/OpenMP will be presented.

# SC99 Panel Discussion: Experiences with Combining OpenMP and MPI

Mary Gabb

A panel of experts will convene at SC99 to address questions concerning hybrid parallelism – i.e., the combination of message passing (e.g., MPI, PVM) and multithreading (e.g., Pthreads, OpenMP). The panelists are experienced in combining distributed- and shared-memory programming methods and have different backgrounds in developing parallel applications: developer/user, the professional developer, the performance analyst, and the end user.

Panelists will give individual presentations and discuss lessons learned in developing and running applications for hybrid parallelism so that future use of SMP clusters, from Beowulf clusters to high-end parallel systems, will be optimized. Panelists will also take questions from the audience.

The panel has been organized, and will be moderated, by Bob Kuhn, Ph.D., of Kuck & Associates, Inc. (KAI). Dr. Kuhn is a recognized expert in parallel computing, and KAI is an established company in parallel tools development. According to Dr. Kuhn, "What I'd like to see happen is more people thinking about both modes of parallelism working together."

The audience is intended to participate by sharing their "war stories" in parallel computing. The panelists can offer a didactic response based on their experiences. As Dr. Kuhn expressed, "We hope to stimulate discussion and reveal new ideas. Hopefully the audience members will see that it's [parallel computing] not so formidable, that it's easier than it first appears."

Panelists will include the following:

Henry Gabb, Ph.D.
Director of Scientific Computing
ERDC MSRC

Howard Scott, Ph.D.
Physicist
Lawrence Livermore National Laboratory

Stef Salvini, Ph.D. HPC Group Leader Numerical Algorithms Group, Inc.

Greg Gaertner
Principal Software Engineer
Compaq Computer Corporation

Danesh Tafti, Ph.D.
Senior Research Scientist and Associate Director
Application Technologies Division

National Center for Supercomputing Applications

John Levesque
Director of the Advanced Computing Technology
Center
IBM Corporation

# Scientific Visualization of the Future: Getting a Grip on Large Data Sets

John E. West

The goal of scientific visualization is to allow an audience to extract information from a collection of data that is then processed intellectually to create understanding. To paraphrase a well-worn quotation, the purpose of visualization is understanding, not pictures. Visualization has been inseparable from scientific endeavor from the outset of formal investigation. As visual animals we have an enormous amount of internal wiring solely for the purpose of acquiring, processing, and (re)acting based on visual input from our environment. A large part of these activities are preconscious, meaning that the processing of visual data does not require us to divert our attention from other tasks — if it did watching a ball game, walking, or driving would be substantially more difficult.


The common perception of many who use visualization to support their research goals is that visualization as a field is stable, with most research simply smoothing the edges of the pioneering work of the 1970s and 1980s. This point of view is testimony to the impact of pioneering visualization and graphics work on nearly all fields of scientific endeavor. It is, however, incorrect. Visualization today is poised to enter an exciting new phase as both the size and expansion rate of data sets continue to increase. As computational, sensor, instrumentation, and algorithmic technologies continue their explosive growth, the visualization research community is faced with an extraordinary set of new challenges. Producers as well as consumers of visualization technology have recognized that we are past the era in which applying bigger hardware to visualization problems provides the substantive technology breakthroughs required for understanding the new breed of very large data sets.

The ability of scientists and engineers to generate or acquire data presently exceeds our ability to process what is gathered in many fields. This represents a gross inefficiency in the information marketplace — we

simply cannot justify not consuming the data that we as a society produce. We are faced with the choice of either scaling back our data initiatives to the level we can process and understand with current technologies or advancing our data exploration abilities to meet new challenges.

Programs like the DoE's Data and Visualization Corridors, the NSF's Large Scientific and Software Data Set Visualization program, and the efforts of many groups akin to the DoD HPCMP's Scientific Visualization Working Group are dedicated to finding innovative hardware/software approaches to the new challenges in visualization. Visualization research is expanding to meet the growing size, distribution, and complexity of data sets while simultaneously addressing the needs of users separated from their data source by both distance and hardware capability. Initiatives in data mining and knowledge discovery, multiresolution data modeling, data storage and access, transparent exchange among heterogeneous databases and disciplines, automatic summarization, and feature extraction are only a few of the broad categories of active research.

The ERDC MSRC is poised to transfer the benefits of this renaissance to our user community. Our visualization infrastructure includes both state-of-the-art hardware and the expertise to apply this hardware with the appropriate software and intellectual capital to facilitate understanding. And our PET partners, outreach, program, and cross-agency collaborations ensure that evolving technologies are identified, researched, and introduced into the infrastructure as they develop.



Dr. Michael Stephens of the Scientific Visualization Center at the ImmersaDesk™

# **More information Customer Assistance Center:**

info-hpc@wes.hpc.mil www.wes.hpc.mil 1-800-500-4722

# interview with . . .

Mary Gabb

According to Mike McCraney, he began his professional life as a "real live" programmer, creating user interfaces for Ada and C simulations for the Army's Space and Strategic Defense Command. He has been at the ERDC (Vicksburg, MS) for 6 years and began his work in the ERDC MSRC Scientific Visualization Center (SVC) approximately 2 years ago as the Group Lead. Mike was the nucleus of the group, organizing the key members of the current ERDC MSRC SVC team. He also worked with the SVC team in acquiring the current impressive hardware and software configurations of the ERDC MSRC SVC. In August, Mike handed control over to Dr. Richard Strelitz, who is now the Acting Scientific Visualization Manager.

Can you tell me a little bit about the history of the ERDC MSRC SciVis [Scientific Visualization] Center? I'd been the Vis [Visualization] Lead for 2 years. At the beginning it was just me. But over the last 2 years, we developed a core team, each [team member] with his own specialty.

We have two visualization specialists with Ph.D.'s. Dr. Kent Eschenberg has spent his entire career doing visualization work for DoD projects. He is also an AVS/Express [Application Visual Systems, Inc.] expert (very innovative in getting around the limitations of AVS/Express and time-varying data sets).

Dr. Richard Strelitz's expertise lies in visualization theory. He came from our PET team, where he was responsible for investigating new visualization technologies and their application.

We also have a video specialist, David Longmire, who has the unique ability to zero in on what the customer needs. You need to get to that 90-percent point before the first review. David is also an expert with the nonlinear [video] editor. We also have an advanced displays specialist, Richard Walters, and an animation specialist, Glen Browning.

Tell me about the Visualization Center's capability. Beginning at PL2 [Performance Level 2], we primarily upgraded memory, disk, and processors in the Center's resources. Since the PL1 graphics hardware satisfied the PL2 requirements, graphics upgrades were not necessary except at the lower end. We introduced three dual-processor SGI Octane systems to satisfy those requirements.

We also started out in PL1 with an ImmersaDesk [Pyramid Systems, Inc.]. One of our Onyxes is dedicated to drive it. However, this machine is also available for normal visualization use when it's not driving the ImmersaDesk. Last year, [the ERDC] ITL [Information Technology Laboratory] purchased a Panoram system,\* which was integrated into an impressive facility known as the Collaboratorium.

Mike McCraney

At PL2, we also acquired a nonlinear editor (Sony) and upgraded to an all-digital router frame in the video suite. All videos are now edited digitally. It's helped tremendously with our efficiency – it allows you to cut, drag and drop all of

allows you to cut, drag, and drop all components (audio, video, transition).

At PL3, we replaced one of the Class III Onyxes with an Onyx 2 with three InfiniteReality2 graphics pipes. That machine drives the Panoram during demonstrations, but users have simultaneous access to the other available graphics pipes. So three users can simultaneously run graphic-intensive jobs. It also has two consoles, so effectively, we have two machines: one can be driving the Panoram system while the other two handle independent graphics jobs.

We also have a dual-processor 540 NT from SGI, one of the new line of visual PCs from SGI. We are in the process of loading as much of the software as we have in the regular Unix environment on it so we can do some benchmark comparisons: AVS/Express, PV-WAVE [Visual Numerics], EnSight [Computational Engineering International], MATLAB [The MathWorks], Data Explorer [IBM]. It will be an interesting matrix when we are done. We can do an "apples-to-apples" comparison, which will be interesting to anyone buying lower end machines.

If I wanted to do a project with ERDC MSRC SciVis Center, how do I set that up? Typically, we have a user who has made some runs on the HPC systems. They realize the importance of visualization and want to work with the Center to gain insight into their research problems and develop professional products to highlight their research. So they look at our web site and see what to do, whom to call. We'll then set up a meeting in the Center or talk to them on the phone or by e-mail about what their requirements are. We discuss their project, what their expectations are, and the context and format of their data sets.

We ask the users to send us sample data sets. With those, we usually do a quick visualization to produce some samples for them to look at. That way, we can see if we are on the right track. In essence, we establish a working relationship where we help them with their visualization or we do it for them. Then we show them what we did so that they can learn to do it themselves the next time around.

<sup>\*</sup> The Panoram system consists of a 30- by 90- foot screen and three Electrohome projectors that can do automatic edge-blending of an image.

Can you give me an overview of the types of projects that you work on? Big projects — such as the RF [Radiofrequency] Weapons, Topside Communications, and the Groundwater Modeling projects — are usually related to Challenge projects. The RF Weapons Challenge project last year is a good example. They were remote users. We found common ground with the RF Weapons people because they had just purchased AVS but didn't yet know how to use it. So that became our collaborative tool. We could send them [AVS] "solution files," and they could use those files to do their own visualization at their own site. We gave them some tools that were really useful, and they continue to use them.

The RF Weapons project is a shining example of how the SciVis Center is supposed to work because we not only provided visualization assistance to a remote user, but we also provided them tools that they continue to use on their own.

We have projects from that scale down to very small projects where the user may be local. The user is here and says, "I really need a video to show my sponsor what I've been doing."

These projects can last from 3 days to a month. Challenge users have the biggest data sets. With those size data sets, they usually won't have local resources that will allow them to do visualization and so they require a center like this.

What are the future plans of SciVis? We're going to continue to investigate new display technologies and new visualization techniques to stay on the cutting edge. That's part of our charter. For example, large data sets are a significant challenge.

Second, we're going to continually investigate collaborative visualization technologies. Collaborative tools are going to be key in the future. We're going to have to support those remote users as if they were sitting here. And we're looking at those technologies now, such that users can log in and actually use the capabilities we have here. While they are logged in, we could theoretically interact with them on-line to help them learn the tools such that they can ultimately do the visualization on their own, without our assistance.

#### Conclusion

As Mike handed the baton to Dr. Strelitz, he reflected on our success to date. "We as an MSRC have focused on the challenge of supporting remote users as well as local users. We have to make the technology and the tools available to as many of our users as we can. That's the whole point of the MSRC scheme, from high performance computing on down."

# interview with

Dr. Richard Strelitz earned his
B.S. in geology of planetary
science at the California
Institute of Technology
just as we were landing
on the moon. His Ph.D.
thesis work was in the area
of seismology. Throughout his

professional development as a scientist over the years, he realized that understanding comes from pictures. According to Dr. Strelitz, that's what the practice of science – seeking the truth – really was. Science is nature visualized. Dr. Strelitz is now the Acting Scientific Visualization Manager at ERDC MSRC.

What are some of the future challenges of scientific visualization in general? I think the future of scientific visualization is wide open because as we

# Richard Strelitz, Ph.D.

become able to deal with larger and more complex problems, we need ways of compressing that data into usable information. And the most obvious thing to do is to put it into pictures, to have the data mapped onto our mental image of what's occurring. So we need to have the science visualized.

The biggest problem that we face is the "disconnect" between the scientists and the visualization. They don't know what's possible with visualization, and we don't always know enough about the science to bring out the details that are useful. One of the key things we've been able to do is to work in partnership with the scientists to find out what they are trying to determine, what they want to explain, what they want to depict, and what they need to show as background, to set the scene. The more complex their models become, the more difficult it is to get a good representation of

the data. But I think there is convergence in understanding. And that's what we're really after.

We've been very successful over the past few years in helping scientists make images and depict representations of their data that not only communicate the content to a lay audience but also mean something to the scientist. And it becomes a feedback process in which the more they see, the more they realize they have to explore.

Can you give an example of a project with a huge data set? Is the scientist struggling to represent it **visually?** Well, perhaps the simplest things are the general climate models, where scientists try to model ocean behavior in response to changes in weather or salinity over a long period of time. The models, in order to be accurate, have to duplicate small details at a very high resolution over the entire ocean, at a variety of depths. Looking at even one 3-D [threedimensional] image is hard enough, but when you want to find out how El Niño affects the weather for 2 years, you need to have several thousand frames of data. How do you maintain a sense of cause-and-effect over a 5-min sequence? How do you know if the boundary conditions are affecting the output in ways that are not intended? This is the problem.

There's a great need for validating computational models and understanding the relationship between data at a given point at a given time, and any other datum. This is the challenge.

Specifically, how do we make more of an effort to communicate with the scientist? The [SciVis] team has come to understand that not all data are equally meaningful, and so we work very hard with the scientists. We don't ask, "Do you want more blue?" but rather, "Is the blue part important to you? What are you trying to look at here? What do you want to focus on? Is it okay if we distort the lens? If we modify the representation, will it beg questions from a lay audience, or will it answer things before they are brought up?" These are the [analytic] questions that our team has been trained to organize. And I really think that from our most-trained scientists to our newest people, we understand the importance of working with the client to come to a mutual understanding of the best representation of the data, while maintaining faithfulness to the science.

Would you say that the process involves more iterative communication with the researcher, rather than "Here's what I need"? Yes, we have very few projects where someone drops off a data set and expects a video by 5:00 P.M. When this happens, the scientist invariably sees the first cut and says, "That's not what I expected to see. I can't make heads or tails of this. What's going on here? Where is my coastline? Where are the reference data that I need in here? I need that swirl, that vortex, that eddy that I didn't see before." Scientific visualization is an interactive process.

We try to have a good initial interview and maintain frequent contact either via e-mail, phone, or face-toface with the scientists. That way, the scientist isn't surprised by what he or she sees later.

As the new Acting SciVis Manager, where would you like to take the SciVis Center? I'd like to think it's going to expand, because if DoD scientists can't look at the data they generate, what's the point of it? Right now, we are limited by what we can see and how we can see things. And as we add greater complexity—either by increasing mesh resolution, or making the mesh more adaptive, shortening the time-step, adding more physical processes—that has to be manifest in our ability to drill down through the data to bring out understanding. And that's the next challenge.

It is not enough to show the electric or magnetic field components generated by a communication array on a ship. To compare behavior at different frequencies or from different antenna configurations, a picture or a clip just won't do. You want to bring that out. You've got, in most situations, an almost incredible degree of complexity, and you want to make use of it. Otherwise, the sponsor will say, for example, "Why did you spend 3,000 node-hours to give me a number that is between 3 and 7? You generated a gigabyte of data. I want to see that gigabyte to see what's useful."

### **Conclusion:**

Scalable parallel computers allow DoD scientists to tackle huge computational problems – with correspondingly large data sets. The mission of the SVC is to enable scientists to interpret their data and, in many cases, see phenomena that they might otherwise miss. Dr. Strelitz's ambitions reflect those of the scientific visualization industry. "We are on the cusp of bringing in new technologies. I think that the DoD recognizes that SciVis is critical to its research objectives. We're at the forefront of it, and we intend to stay there."

# A Helping Hand in the Customer Assistance Center

Robin Phillips

The Customer Assistance Center (CAC) is the user's advocate at the ERDC MSRC. It is responsible for technical support and account administration. A rotating group of six "frontline" support analysts staff the CAC from 0600 to 1900 CDT, 5 days a week (operators answer the phones during the "off" hours). The analysts receive e-mail, telephone, and walk-in requests for assistance with a priority of responding quickly and accurately to the users. The CAC members listen to the users' needs on an individual basis, solve first-level problems if possible, and, if necessary, convey those needs to the appropriate expertise in the MSRC.

CAC personnel enter requests into a call-ticket database and track the problem through the database until resolution. The originating CAC member assumes ownership of the problem. The database allows each CAC member to always have a clear understanding of a problem and its progress to closure, irrespective of which CAC member owns the problem.

The CAC uses a layered support approach. Most problems and questions are addressed directly by CAC personnel, while others are forwarded to a second layer of expertise, the application analysts. Requests that require the attention of system administrators are forwarded to the system administrator(s) for that particular system. Hardware and third-party software problems that cannot be resolved by systems or applications personnel are forwarded to the third layer, which consists of onsite vendor representatives, specialists from the PET program, the Computational Migration Group, or the Computational Science and Engineering Group. This layered support approach is effective in providing users with the proper level of expertise in a timely manner.

The CAC has recently implemented a new web- and Oracle-based Administration, Allocation, and Utilization (AAU) system for maintaining and reporting user and project information. The AAU system contains the database, the user interface, administrative processes, and reporting mechanisms necessary for maintaining and supporting the ERDC MSRC's 223 projects, 884 users, and 2,296 user accounts.

A customer satisfaction survey is conducted quarterly to seek out customers' opinions and problems. Many process improvements can be directly traced to the survey responses. The surveys consistently reveal that more than 90 percent of respondents rate the ERDC MSRC CAC as "excellent" or "good" in courtesy, knowledge, and total customer support quality. Suggestions and feedback are always welcome (please see "More Information" on the next page).

The goal of the CAC is to evaluate every problem and provide a solution in the shortest time available. The challenge is to meet this goal in the dynamic ERDC MSRC environment.

# Looking at CAC Through a 1-Year Window

Total Number of Service Calls	4,380	
<b>Number of Users</b>	884	
<b>Number of User Accounts</b>	2,296	

### Lifespan of Service Call

<ul><li>Average</li></ul>	<b>4.22 days</b>
	(101.21 hr)
<ul><li>Mode</li></ul>	0 days (0 hr)
<ul><li>Median</li></ul>	0.08 days
	(1.87  hr)

These data represent the activities of the ERDC MSRC CAC over FY99. The average lifespan of a service call is about 4 days. However, most service calls are resolved within an hour.



Six "frontline" support analysts staff the CAC from 0600 to 1900 CDT daily during standard government hours. Shown left to right, front row: Lisa Langford, Melissa Hampton, Kelly Lanier. Left to right, back row: Mike Gough (CAC Manager), Frank Green, Robin Phillips, Jay Sykes.

# **More information Customer Assistance Center:**

info-hpc@wes.hpc.mil www.wes.hpc.mil 1-800-500-4722

### **PET Profiles**

Mary Gabb

### Clay Breshears, Ph.D.

The Rice University Onsite Lead for Scalable Parallel Programming (SPP) Tools

Dr. Clay Breshears, a native of Spokane, WA, has been in Vicksburg as part of the ERDC MSRC PET program since June 1997. As the onsite Scalable Parallel Programming Tools Lead, he is employed as a research scientist by the Center for Research on Parallel Computation at Rice University. In this position, Dr. Breshears determines the computational and programming support needs of ERDC MSRC users, evaluates available software and systems, recommends new programming tools, and helps users get the most out of available tools. He also occasionally collaborates with Rice University researchers to find solutions to ERDC MSRC tools-related issues or problems.

Dr. Breshears received his B.S. in computer science from Eastern Washington University (EWU), his M.S. in computer science from Washington State University, and his Ph.D. in computer science from the University of Tennessee at Knoxville. Dr. Breshears' research interests include parallel and distributed computation, concurrent programming, and space-optimal and non-numeric parallel algorithms. Dr. Breshears also follows other academic interests

including theoretical computer science, genetic algorithms, cryptography, and methods for using DNA molecules as computational devices.

Most of Dr. Breshears' professional career has been spent in university settings. He was on the computer science faculty for 4 years at EWU (where he also served 1 year as the Department Chair) and 2 years at the University of Southern Mississippi. He has

taught workshops in high performance computing, multithreaded and concurrent programming, and various parallel programming tools. Dr. Breshears has been involved with the parallel conversion of a variety of scientific programs including LAPACK routines to create an eigensolver for complex, symmetric, non-Hermitian matrices.

Outside of work, his leisure time is spent on other cerebral activities. He enjoys theater and movies, Japanese manga and anime, watching the Chicago Cubs, and playing chess (U.S. Chess Federation member). He is also a fierce Mah Jongg competitor and collects comic books (since 1970). Dr. Breshears also does aerobics, lifts weights, and plays softball to stay active.

### Wayne Mastin, Ph.D.

Onsite PET Academic Team Lead, Nichols Research Corporation Professor Emeritus, Mississippi State University (MSU)

Professor Wayne Mastin has had a fruitful and highly decorated career in mathematics and high performance computing. He received his B.S. in mathematics from Austin Peay State College in 1964, his M.S. in mathematics from Miami University in 1966, and his Ph.D. in mathematics from Texas Christian University in 1969.

At the ERDC MSRC, which he joined in 1996 as the onsite Academic Team Lead for PET, he oversees the training program and provides senior leadership for other onsite PET leads. Since receiving his Ph.D., he has held several positions with the National Aeronautics and Space Administration (NASA) and has served as a professor in the Department of Mathematics and

Statistics at MSU. He was appointed Professor Emeritus of Mathematics and Statistics at MSU in 1995. His 25 years of teaching experience have been with MSU, including courses at the ERDC Graduate Institute in Vicksburg, MS. He has co-authored a book on grid generation and authored numerous articles.

Professor Mastin has particular interests in grid generation, numerical analysis, numerical methods and software, numerical solutions of partial differential equations, computational fluid dynamics, and computational geometry. Throughout his career, he has been involved in the development and application of a variety of surface modeling, grid generation, computational fluid dynamics, hydrodynamics, and scientific visualization codes.

In his precious free time, Professor Mastin is a tinkerer – on his 1970 Opel GT.

### Richard A. Weed, Ph.D.

The Mississippi State University (MSU) Onsite Lead for Computational Structural Mechanics (CSM)

Dr. Richard Weed received his B.S. in aerospace engineering from MSU in 1974, his M.S. in aerospace engineering from MSU in 1980, and his Ph.D. in aerospace engineering from the Georgia Institute of Technology in 1995.

Dr. Weed is currently employed by the National Science Foundation Engineering Research Center for Computational Field Simulation at MSU. As the ERDC PET CSM Lead, Dr. Weed's activities include providing support for ongoing CSM DoD Challenge projects at the ERDC MSRC, the development of new structural analysis algorithms and computer codes on the ERDC MSRC parallel computing platforms, user outreach, and education in new CSM analysis methods. Dr. Weed also provides programmatic support for the ongoing PET core support and Focused Effort activities of the Texas Institute for Computational and Applied Mathematics (TICAM) at The University of Texas at Austin.

TICAM serves as the PET Academic Lead institution for Computational Structural Dynamics.

Dr. Weed's work experience includes over 15 years as an aerospace engineer with the Lockheed-Martin Aeronautical Systems Company, the Lockheed-Martin Missiles and Space Company, and the

McDonnell-Douglas Engineering Services Company. During his years with Lockheed, Dr. Weed was involved in the development of advanced CFD methods for the design and analysis of a wide range of aircraft, missile, and space vehicle configurations. At McDonnell-Douglas, Dr. Weed worked in the development and validation of space shuttle guidance and control software.

Dr. Weed's research interests include development of efficient multidisciplinary solution algorithms for both distributed- and shared-memory parallel systems, parallel algorithms for fluid-structures interactions, computational methods for explosive effects prediction and analysis, and finite element methods for structural dynamics analysis.

### Phu Luong, Ph.D.

The University of Texas at Austin Onsite Lead for Environmental Quality Modeling and Simulation (EQM)

Dr. Phu Luong has a special relationship with Mississippi. He arrived here in 1981, sponsored by the Trinity Presbyterian Church in Starkville, MS, after leaving his home country, South Vietnam. After immigrating to the United States, he received his B.S. and M.S. in applied mathematics and computer science from the University of Mississippi in 1986 and MSU in 1988, respectively. In 1991, he was also awarded a doctoral degree in computational engineering from MSU.

His postdoctoral career also claims Mississippi residence. He worked at the Naval Oceanographic Office MSRC (Stennis Space Center, MS) from 1991 to 1999, focusing on ocean circulation modeling and

multiblock grid generation applied to coastal ocean modeling.

Dr. Luong's research interests include adaptive grid generation and numerical solution of partial differential equations. As the ERDC EQM onsite Lead (and research scientist for the TICAM at The University of Texas at Austin), he is working with hydraulics and sediment transport models, which have potential application to groundwater contamination studies. He

is also interested in multiphysics phenomena modeling, parallelization, and scientific visualization.

Outside of the office his interests are quite simple – he enjoys evening walks with his family and spending time with his 2-year-old son, Minh.

### John Eberle

PET Training Coordinator

Mr. John Eberle devoted the first 23 years of his adult life to the U.S. Marine Corps, serving several tours in southeast Asia. He assumed increasingly demanding leadership roles, retiring as a Master Gunnery Sergeant in 1987.

Training comprises the vast majority of his professional military work. He created, updated, and/or refined full plans of instruction for the Army, Air Force, the DoE, and numerous Marine Corps courses. He is also a world expert on the planning and operational use of unattended remote ground sensors, having been involved in the design, production, testing, and implementation of unattended sensors since the early 1970s. Mr. Eberle has acted as the Test Director on several Marine Corps Development efforts and worked in the Joint Testing Arena after retirement from the Marine Corps.

As the Training Coordinator for the ERDC MSRC, he is responsible for scheduling and maintaining the training facility and ensuring the smooth and timely implementation of the myriad HPC classes offered at the ERDC MSRC. In addition, he is the onsite leader for the web-based distance training collaborative system (TANGO

Interactive) designed for the ERDC MSRC by Syracuse University. This system is designed to decrease the need for onsite classroom seats by allowing remote participants to "join" an ongoing class without leaving their worksites.

Mr. Eberle is a Microsoft Certified Professional and has earned Microsoft Certified Systems Engineer Certification.

In his spare time, Mr. Eberle is an avid woodworker and reader.

### Stephen Wornom, Ph.D.

The Ohio State University Onsite Lead for Climate, Weather, and Ocean Modeling and Simulation

Dr. Stephen Wornom began his scientific career in the foothills of the Blue Ridge mountains, attending Virginia Tech to earn his B.S. (1967) and Ph.D. (1971) in Engineering Mechanics (a cross-curriculum of mathematics and physics).

His postdoctoral experience involved developing and applying numerical methods to compressible aerodynamic flows at the NASA Langley Research Center.

Later, Dr. Wornom seized the opportunity to work in France, an adventure that would last 7 years. At PRINCIPIA, an ocean engineering company, Dr. Wornom applied his CFD experience to hydrodynamic flows. His group developed new pseudocompressibility formulations to solve the

incompressible Navier-Stokes equations so that numerical methods developed for solving the compressible equations could be directly applied to solve the incompressible equations.

In 1997, he returned to the United States

to work on the hypersonic SCRAMJET program (U.S. Air Force) at Wright-Patterson Air Force Base. He developed an application for designing hypersonic inlets and performed CFD flight simulations for SCRAMJET designs.

Dr. Wornom's current scientific interest is the coupling of nearshore wave models with deep-ocean models to better predict nearshore behavior.

Dr. Wornom's outside interests include game theory. He also walks several miles daily and plays classical guitar and nonclassical piano.

### Henry A. Gabb, Ph.D.

PET Director, Director of Scientific Computing Nichols Research Corporation

Dr. Henry A. Gabb's background in biochemistry and scientific computation gives him a unique approach to handling the scientific problems commonly addressed in high performance computing. He received his B.S. in biochemistry from Louisiana State University. Molecular biologists were just starting to use computers for genetic sequence analysis, so this was also his first exposure to scientific computation.

He was awarded a Ph.D. in biochemistry and molecular genetics from the University of Alabama at Birmingham Schools of Medicine and Dentistry. His thesis used computational methods to study conformational transitions in nucleotides, the building blocks of DNA. Dr. Gabb received a postdoctoral fellowship from the French Foreign Ministry to do research at l'Institut de Biologie et Physico-Chimique in Paris. There he used Monte Carlo simulation techniques to study DNA flexibility. He also used computational chemistry to study base-stacking interactions in DNA.

After his postdoctoral study, Dr. Gabb moved to the Imperial Cancer Research Fund in London, where he applied Fourier correlation techniques to computer-aided molecular docking.

Dr. Gabb is currently the Director of Scientific Computing at the ERDC MSRC. He is the author of more than 20 research articles and invited review

articles on a wide range of scientific topics including molecular modeling, human genetics, and high performance computing. He is a Fellow of the Imperial Cancer Research Fund and a Fellow of the European Molecular Biology Organization. His scientific interests include multithreaded programming, SMP cluster architectures, the C3I Defense Systems Benchmarks, and computational biology.

In his spare time, Dr. Gabb enjoys learning new programming languages and, if you can get him away from the computer, studying Japanese history and culture, and brewing his own beer.

# **ERDC MSRC Training**

Wayne Mastin, Ph.D.

The ERDC MSRC PET program is dedicated to the task of keeping our users abreast of the latest developments in HPC technology. This includes training on how to access and efficiently use our HPC hardware and how to use new programming methods. The PET program offers training classes supported by the PET university team and our onsite leads, as well as classes offered by hardware and software vendors. The PET training team also stands ready to assist with usersponsored training events and workshops. The ERDC MSRC Training and Education Facility contains equipment to project, record, and broadcast over the Internet training material in nearly any format chosen by the presenter.

A major emphasis of training at the ERDC MSRC continues to be, as it has from the beginning, the delivery of information to the remote user. The vast majority of ERDC MSRC users are located at remote sites around the United States. Over the last 3 years, we have progressed from satellite TV and MBONE broadcasts to the richer web-based Tango Interactive

collaborative system. The requirements for receiving Tango at remote locations are minimal, and the PET training staff is available to assist users in setting up their computer systems to participate in training classes.

The development of Tango as a tool for distance training has been a joint effort of several PET partners. Syracuse University originally developed the Tango Interactive technology in cooperation with the Air Force Research Laboratory-Rome Research Site as a tool for collaboration over the web. The further development of Tango as a tool for distance education and training was a joint effort of Syracuse and Jackson State University (JSU). JSU is now in its third year of offering full-semester, college-credit courses using Tango and delivered a course this past year to Morgan State University. HPC training material for Tangobased classes has been developed by Syracuse and the Ohio Supercomputer Center with PET support.

While methods for the delivery of training to remote sites have improved, some offsite users may still prefer the intimacy of a classroom environment. For those users, we extend an open invitation for them to attend any of our onsite training classes. Information and assistance are available on directions, lodging, and leisure activities to fill their free time. In cases where a sufficient number of interested participants are all from one site, we also offer training classes at remote sites.

The PET training program is continually evolving as new technology and new hardware enter the market place. In the beginning, the primary emphasis was on techniques in parallel programming. The aim was to assist users in understanding programming for scalable parallel systems and to get them to use these systems as quickly as possible. Topics included parallel methods, such as Message-Passing Interface, OpenMP, and High Performance Fortran, as well as debuggers, performance profilers, and parallel libraries (such as ScaLAPACK and PETSc). These topics are still important, and many are covered in our hardware-specific courses. However, the emphasis in training is now more domain specific. Courses are targeted at particular HPC software or specific Computational Technology Areas. We have moved from an objective of getting the user on the parallel systems and running a few trial examples to assisting the user in implementing the software or technology required to address large scalable problems. Recently, workshops were held on computational methods in grid generation and CWO, and training classes were conducted on the Dynamic Analysis and Design System (for mechanical system analysis) and EnSight (for scientific visualization) software packages. The PET team also collaborated with the Defense Modeling and Simulation Office to offer training in highlevel architecture standards for interoperability between various DoD modeling and simulation systems. In addition, the arrival of the next addition to the ERDC MSRC stable of HPC systems, the IBM Power3 SMP, will be supported by a training class to apprise users of the expanded capabilities and nuances of the new architecture. Although the topics covered by the PET training program have changed, the primary mission remains unchanged. That mission is to support the DoD research and development user community. The arrival of the next addition to the ERDC MSRC stable of HPC machines, the IBM Power3 SMP, will be followed by a training class to apprise users of the expanded capabilities and nuances of the new machine.

Because most of the instructors have a continuing relationship with the PET program, quite often knowledge exchange does not end on the last day of class. Training classes have led to extended communication and collaboration on research projects between instructors and participants and among individual participants. In the majority of cases, the classes offered under the PET program are focused on a specific topic and tend to draw participants with similar or complementary interests, and collaboration is a natural outcome. Some of these collaborations have resulted in university projects sponsored by the PET program.

The PET team solicits suggestions for new training activities. One of the missions of the PET team is to match the needs of ERDC MSRC users with the capabilities of our onsite staff, the university team members, and training available from our software and hardware vendors. Suggestions may be sent to our Customer Assistance Center (*info-hpc@wes.hpc.mil*, or 1-800-500-4722).



The ERDC MSRC Training and Education Facility contains equipment to project, record, and broadcast over the Internet training material in nearly any format chosen by the presenter

### **Training Made to Order**

### Henry A. Gabb, Ph.D.

A popular course offered by ERDC MSRC is the Bring-Your-Own-Code Workshop taught jointly by the Computational Migration Group and the PET team. As the name implies, attendees are expected to bring a code they would like to migrate to or optimize on a parallel system. This workshop begins with a 1-day overview of parallel programming methods and parallel architectures. The attendees spend the next 3 days parallelizing their codes with direct assistance from the CMG and PET teams. This course requires a training facility equipped with enough workstations for each attendee.

BBN Technologies (under contract to Navy) recently requested a parallel programming workshop on behalf of ERDC MSRC users at nearby Electric Boat Corporation (New London, CT). Enough people registered for the course, so the instructors traveled to the students. The attendees were experienced HPC programmers who were mainly interested in learning new techniques and advanced topics in parallel computing, so it was decided that a lecture format was more appropriate. After consulting the attendees, a course syllabus was produced that consisted of a 1-day lecture on OpenMP, 1 day for MPI, and 1 day for advanced OpenMP, performance issues specific to the SGI Origin2000, and issues combining MPI and OpenMP in the same application.

Popular courses offered by ERDC MSRC staff include the following:

- Bring-Your-Own-Code Workshop
- ♦ Parallel Programming with OpenMP
- ♥ Parallel Programming with MPI
- Street Concurrent Programming with Pthreads.

By prior arrangement, the courses listed can be tailored to the specific needs of the attendees. If there is sufficient interest in a particular topic, new courses will be produced.

More information
Customer Assistance Center:

info-hpc@wes.hpc.mil www.wes.hpc.mil 1-800-500-4722

# technical reports

- 99-04 Mary F. Wheeler, Clint Dawson, Jichun Li, and Victor Parr, "UTPROJ: The University of Texas Projection Code for Computing Locally Conservative Velocity Fields"
- 99-05 Randy Heiland and M. Pauline Baker, "Coprocessing: Experience with CUMULVS and pV3"
- 99-06 Shirley Browne, George Ho, and Phil Mucci, "PAPI: Portable Interface to Hardware Performance Counters"
- 99-07 Alan M. Shih and M. Polly Baker, "Management Strategies for Scientific Data: Assessing Utility of HDF for CEWES MSRC Users"
- 99-08 Wayne Mastin, "1998 CEWES MSRC PET Training Activities"
- 99-09 J. B. White III and S. W. Bova, "Where's the Overlap? An Analysis of Popular MPI Implementation"
- 99-10 Dan Nagle, "The Effect of Fortran 95 PURE and ELEMENTAL Procedures on Parallel Execution"
- 99-11 Clay P. Breshears and Graham Fagg, "A Computation Allocation Model for Distributed Computing Under MPI\_Connect"
- 99-12 Mark R. Fahey and Dan Nagle, "Cray Fortran Pointers vs. Fortran 90 Pointers and Porting from the CRAY C90 to the SGI Origin2000"
- 99-13 Dan Nagle, "The Benefits of Fortran Procedure Interfaces"
- 99-14 Geoffrey C. Fox, Wojtek Furmanski, Subhash Nair, Hasan T. Ozdemir, Zeynep Ozdemir, and Tom A. Pulikal, "WebHLA An Interactive Multiplayer Environment for High Performance Distributed Modeling and Simulation"
- 99-15 Geoffrey C. Fox, Wojtek Furmanski, Ganesh Krishnamurthy, Hasan T. Ozdemir, Zeynep Odcikin-Ozdemir, Tom A. Pulikal, Krishnan Rangarajan, and Ankur Sood, "Using WebHLA to Integrate HPC FMS Modules with Web/Commodity based Distributed Object Technologies of CORBA, Java, COM and XML"
- 99-16 Tomasz Haupt, Erol Akarsu, and Geoffrey C. Fox, "Landscape Management System: A WebFlow Application"
- 99-17 Benjamin Willhoite and Dan Nagle, "Using Fortran 90 Features for Cache Optimization"
- 99-18 Steve W. Bova, Clay P. Breshears, Christine Cuicchi, Zeki Demirbilek, and Henry A. Gabb, "Dual-level Parallel Analysis of Harbor Wave Response Using MPI and OpenMP"
- 99-19 Sandie Kappes, "PET Web Pages Evolution"
- 99-20 O. Olatidoye, G. Jones, S. Sarathy, C. McIntyre, and L. Milligan, "Scientific Visualization for Interpretation and Assessment of Damage in Structures Subject to Blast Loads"
- 99-21 Troy Baer, David Ennis, James Giuliani, Leslie Southern, and David E. Bernholdt, "Experiences with Using Tango Interactive in a Distributed Workshop"
- 99-22 D. J. S. Welsh, R. Wang, P. Sadayappan, and K. W. Bedford, "Coupling of Marine Circulation and Wind-Wave Models on Parallel Platforms"
- 99-23 P. V. Bangalore, J. Zhu, D. Huddleston, A. Skjellum, D. J. S. Welsh, K. W. Bedford, R. Wang, and P. Sadayappan, "Parallelization of a Coupled Hydraulics and Sediment Transport Model"
- 99-24 Contract Year Four PET Core Support and Focused Efforts for CEWES MSRC
- 99-25 M. Ehtesham Hayder and John Mellor-Crummey, "Improving Performance with Scalar Replacement"
- 99-26 James B. White III, "Reading Sequential Unformatted CRAY C90 Files on an SGI Origin"
- 99-27 M. Ehtesham Hayder, Constantinos S. Ierotheou, and David E. Keys, "Three Parallel Programming Paradigms: Comparisons on an Archetypal PDE Computation"
- 99-28 Chandrajit Bajaj, Steven Cutchin, and Mary Wheeler, "Simulation Code Launching from the Web"
- 99-29 Michael A. Chupa and Robert J. Moorhead, "Oceanographic Model Visualization with the Interactive Structured Time-Varying Visualizer"
- 99-30 Dan Nagle and Joy Brogdon, "Fortran 90 Namelist I/O versus Cray Namelist I/O"

# technical reports

- 99-31 Abani K. Patra, Atanas I. Pehlivanov, David Littlefield, Graham F. Carey, and J. Tinsley Oden, "Application of Error Indicators and Local Adaptive Refinement for Elasto-Plasic Impact Calculation (EPIC)"
- 99-32 Geoffrey C. Fox, Roman Markowski, Nancy J. McCracken, Marek Podgorney, Qutaibah Malluhi, and Debasis Mitra, "More Experiences with Tango Interactive in Synchronous Distance Learning Courses"
- 99-33 Shirley Browne, Paul McMahan, and Scott Wells, "Repository in a Box Toolkit for Software and Resource Sharing"
- 99-34 Rob Stein, "CBayVisGen User Guide"

These technical reports can be accessed at www.wes.hpc.mil/.

## **ERDC MSRC**

# training schedule\*

### December 1999

Using the SGI Origin2000 for Code Development and Analysis

#### January 2000

Coupling Multiphysics Problems in Environmental Simulations

#### January 2000

Workshop on Mesoscale Atmospheric Models

### February 2000

Adaptive Meshes for CSM

#### March 2000

Tango "Open House"

\* Additional courses may be offered. Please check the ERDC MSRC web page at <a href="http://www.wes.hpc.mil/">http://www.wes.hpc.mil/</a>.

# **ERDC MSRC Newsletter Editoral Staff**

### **Interdisciplinary Coordinator**

Mary L. Hampton

### **Outreach Coordinator**

Lee Byrne

### **Chief Editors**

Rose Dykes Mary Gabb

### **Visual Information Specialist**

Betty Watson

### **ERDC MSRC Customer Assistance Center**

Web site: http://www.wes.hpc.mil/ E-mail: info-hpc@wes.hpc.mil Telephone: 1-800-500-4722

Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the DoD. Your comments, ideas, and contributions are welcome.

Design and layout provided by the Visual Production Center, Information Technology Laboratory,

U.S. Army Engineer Research and Development Center.

Approved for public release; distribution is unlimited.

